UNIVERSITY OF
ILLINOIS LIBRARY
AT URBANA CHAMPAIGN
BOOKSTACKS

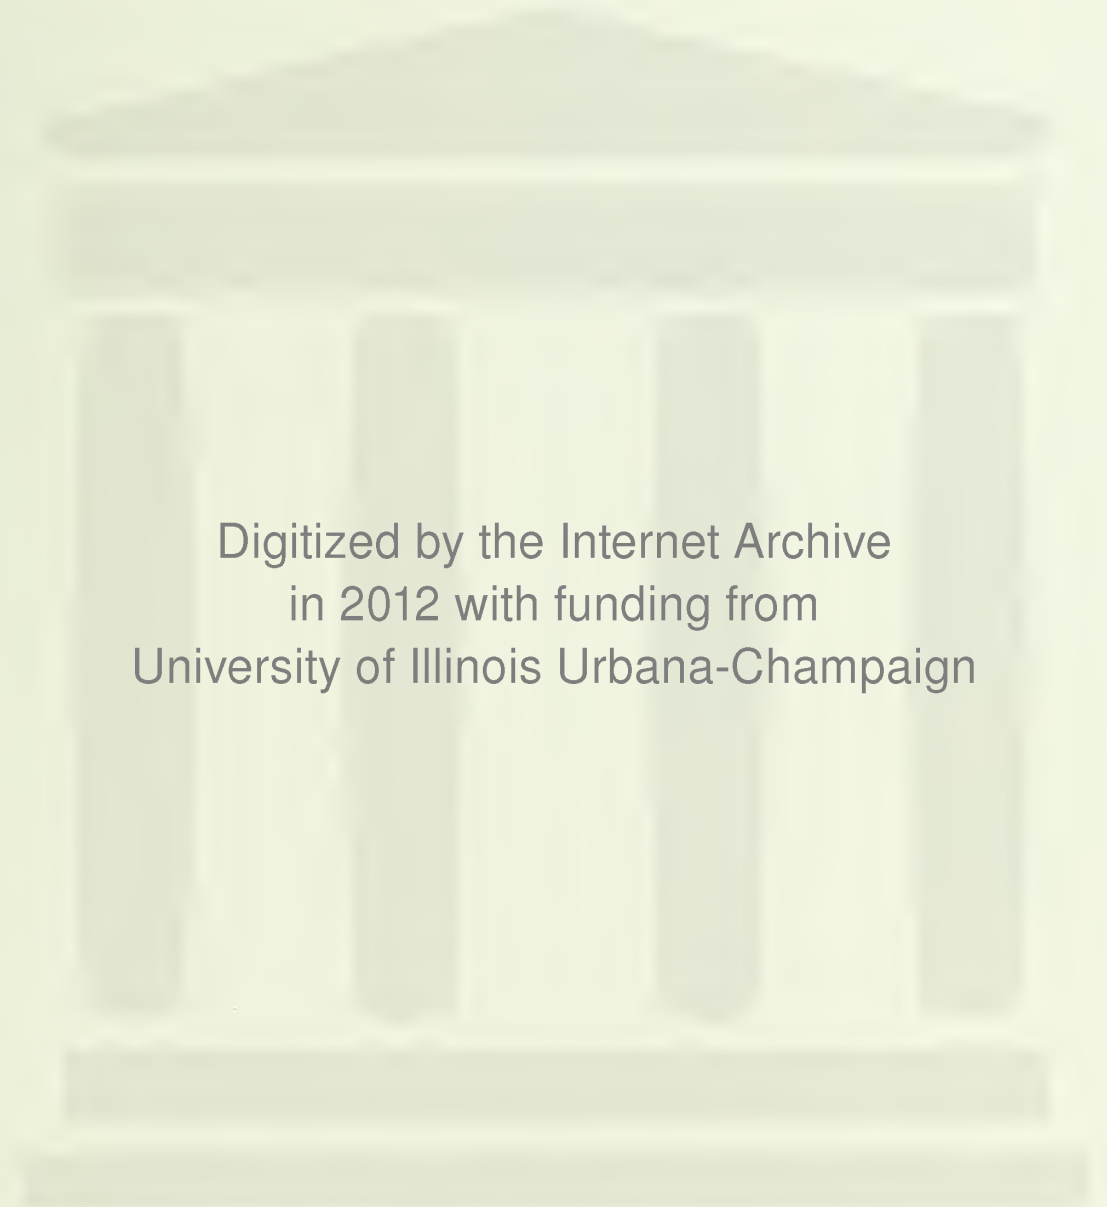

Digitized by the Internet Archive
in 2012 with funding from
University of Illinois Urbana-Champaign

http://www.archive.org/details/cheapplaywithnor91122chak

FACULTY WORKING PAPER 91-0122

STX
330
B385
1991:122 COPY 2

Cheap Play With No Regret

The Library of the
MAY 1 6 1991
University of Illinois
of Urbana-Champaign

Bhaskar Chakravorti
Department of Economics

John P. Conley
Department of Economics

Bureau of Economic and Business Research
College of Commerce and Business Administration
University of Illinois at Urbana-Champaign

BEBR

FACULTY WORKING PAPER NO. 91-0122

College of Commerce and Business Administration

University of Illinois at Urbana-Champaign

March 1991

Cheap Play With No Regret

Bhaskar Chakravorti

and

John P. Conley

Department of Economics
University of Illinois at Urbana-Champaign
Champaign, IL 61820

The authors would like to thank the participants of the economic theory workshop at the University of Illinois, particularly Mark Feldman, for comments and stimulating discussions. The usual disclaimer applies.

CHEAP PLAY WITH NO REGRET

by

Bhaskar Chakravorti

and

John P. Conley

ABSTRACT

This paper studies a special class of differential information games with pre-play communication -- games with "cheap play." We consider problems in which there are several rounds of payoff-irrelevant publicly observable choice (or discussion) of actions, followed by a final round in which actions are binding and payoff relevant. A natural focal subset of equilibria of such games is one that consists of equilibria involving *no regret*. In such equilibria, actions chosen in the cheap play phase remain optimal even after private information is updated and are chosen again at the binding stage. We show that there are several alternative ways of formalizing the notion of sequential equilibria with no regret and present an argument for selecting one such formulation. The arguments rely on the interpretation of mixed strategies and on the specification of regret-freeness off the equilibrium path. We provide a complete characterization of the set of regret-free sequential equilibrium outcomes of the extended game in terms of a "posterior implementability" criterion applied to the underlying static game.

JEL Classification: 026

Address correspondence to:

John P. Conley
Department of Economics
University of Illinois at Urbana-Champaign
330 Commerce West
1206 South Sixth Street
Champaign, IL 61820

1. INTRODUCTION

The focus of this paper is on an important sub-class of games with pre-play communication -- games with *cheap play*. We consider problems in which there are several rounds of payoff-irrelevant publicly observable choice (or discussion) of actions, followed by a final round in which actions are binding and payoff relevant. If there is differential information among the players, such "cheap play" in the earlier rounds conveys information. The crucial aspect of cheap play that sets it apart from other forms of pre-play communication is that natural focal equilibria of the resulting game are those involving *no regret*. In such equilibria, actions chosen in the cheap play phase remain optimal even after private information is updated and are chosen again when the time comes to make a binding commitment.

We develop a formal model of the sequential game induced by the imposition of rounds of cheap play to an underlying simultaneous-move game. The objective is to identify the appropriate formalization of equilibria involving no regret. We present seven alternative candidates. These involve various types of stationarity restrictions on the sequential equilibria of the induced game. We argue that other definitions will not differ in any essential way. We make the case for singling one of them out and provide a simple characterization of the set of sequential equilibrium outcomes satisfying this form of stationarity.

The definition of the set of equilibrium outcomes of a game with pre-play communication via non-binding play is potentially complicated. However, our main theorem shows that regret-free equilibrium outcomes may be characterized by a set of linear inequalities. The practical

implication of this is that the design of optimal contracts in situations where such pre-play communication occurs becomes a straightforward programming problem.

The central theoretical implication of our paper is the establishment of a link between the appropriate formulation of no regret and an equilibrium refinement proposed by Green and Laffont (1987). Green and Laffont's objective is precisely the same as ours -- to study the effect of pre-play rounds of cheap play with no regret. They do not explicitly analyze the sequential game induced by such pre-play communication; instead they impose a restriction called *posterior implementability* on the Bayesian equilibria of the underlying static game. However, Green and Laffont speculate that there may be some connection between the perfect equilibria of the extended game induced by pre-play communication a la Farrell (1982) and posterior implementability.

Our paper addresses this speculation. We demonstrate an equivalence between the set of stationary sequential equilibrium outcomes of the game with cheap play and the set of posterior implementable Bayesian equilibrium outcomes of the underlying static game. Hence, this paper can be interpreted as a justification for directly concentrating on the static game, as Green and Laffont did, without bothering with the messy extended sequential game. In the absence of such an equivalence argument, an analysis that rests on a restriction on the Bayesian equilibria of the static game would be rather troubling.

Before we plunge into the formalities, we shall briefly remark on the game-theoretic implications of the various formulations of no regret. Of the seven possible alternatives, we show that four (which involve regret-freeness off the equilibrium path) are too strong and are therefore

trivial. Of the remaining possibilites, the issue, quite unexpectedly, reduces to the interpretation that one adopts with regard to mixed strategies. The "modern" view of Harsanyi (1973) and Aumann (1987) is that mixed strategies should be treated as an expression of the uncertainty that other players have about the strategy choice of any given player. This is in contrast to the "traditional" view that mixed strategies involve players making decisions on the basis of actual "flips of a coin". We show that posterior implementability completely characterizes the set of sequential equilibrium outcomes satisfying a notion of regret-freeness consistent with the Harsanyi-Aumann viewpoint. The candidates that are consistent with the traditional viewpoint have no logical relationship with posterior implementability.

Section 2 contains the basic model. Section 3 presents alternative ways of formalizing the notion of no regret. Section 4 contains the results and the final section concludes.

2. THE MODEL

A is a finite set of *outcomes*, $N = \{1, 2, \ldots, n\}$ is a set of *players*, and Θ is a finite set of *states of the world*. A state of the world $\theta \in \Theta$ is a profile $(\theta_i)_{i \in N}$. Each player i in N is characterized by a list $\langle \Theta_i, u_i, \tilde{\pi}_i, M_i \rangle$, which includes:

a set of possible *private observations* Θ_i,

a (von Neumann-Morgenstern) *utility function* u_i: $A \times \Theta \to \mathbb{R}$,

a *prior probability distribution on* Θ, $\tilde{\pi}_i$: $\Theta \to (0, 1]$, and

a finite set of *moves*, M_i.

In the sequel, for any set X_i, let $X \equiv \times_{i \in N} X_i$, $X_{-i} \equiv \times_{j \in N \setminus (i)} X_j$ and given $x_i \in X_i$, let $x \equiv (x_i)_{i \in N}$ and $x_{-i} \equiv (x_j)_{j \in N \setminus (i)}$. Also, for any set Y, let $\Delta(Y)$ denote the set of randomizations over Y. Given a random variable h: $X \to \Delta(Y)$, we shall, with some abuse of notation, use $h(y \mid x)$ to denote the probability assigned to $y \in Y$ by the distribution $h(x)$.

An *outcome function* g: $M \to A$ specifies an outcome for every profile of moves. A *game (form)* Γ is a list $\langle A, N, M, g \rangle$. The *strategy space* for i in the game Γ is the set $S_i = \{s_i: \Theta_i \to \Delta(M_i)\}$. With slight abuse of notation, $s(\cdot \mid \theta)$ and $s_{-i}(\cdot \mid \theta_{-i})$ denote the joint probability distributions induced on M and M_{-i} by $s \in S$, and $s_{-i} \in S_{-i}$, respectively, given a realization $\theta \in \Theta$. Let the function g^*s: $\Theta \to \Delta(A)$ be defined by: for all $a \in A$ and $\theta \in \Theta$,

$$g^*s(a \mid \theta) = \begin{cases} \sum_{m' \in (m:\ g(m) = a)} s(m' \mid \theta) & \text{if } \exists m \text{ such that } g(m) = a \\ 0 & \text{otherwise.} \end{cases}$$

This is the function that specifies the probability distribution induced on the set of outcomes given the choice of s in the game played in the state of the world θ.

Every player i updates his/her probability distribution on Θ upon observing an element of Θ_i using Bayes' Law. This is summarized by a *posterior probability distribution on* Θ_{-i}, π_i: $\Theta \to [0, 1]$, where $\pi_i(\theta)$ specifies the probability assigned by i to $\theta_{-i} \in \Theta_{-i}$, given the observation $\theta_i \in \Theta_i$.

The model thus far is assumed to be common knowledge in the sense of Aumann (1976).

A *Bayesian equilibrium* of Γ is a strategy profile $s \in S$ that

satisfies:

$$\forall i \in N,\ \forall \theta_i \in \Theta_i,\ \forall s_i' \in S_i,\ s' \equiv (s_i', s_{-i}),$$

$$\sum_{\theta_{-i} \in \Theta_{-i}} \sum_{a \in A} \pi_i(\theta) g^* s(a \mid \theta) u_i(a, \theta) \geq \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{a \in A} \pi_i(\theta) g^* s'(a \mid \theta) u_i(a, \theta)$$

$\mathcal{E}_S(\Gamma) \subseteq S$ denotes the set of Bayesian equilibria of Γ, with $\mathcal{E}(\Gamma) \equiv \{g^* s: s \in \mathcal{E}_S(\Gamma)\}$ denoting the set of induced Bayesian equilibrium outcomes of Γ.

Next, we shall model the following situation. Prior to committing themselves to moves in the game, the players engage in pre-play communication. The communication is in the form of "cheap play" -- each player makes or proposes a move for himself/herself with the knowledge that it can be withdrawn at no cost. We allow for several rounds of such non-binding and payoff-irrelevant (simultaneous-move) play.

Let T denote the number of rounds of simultaneous-move choices available to the players such that the choices made from M in the first $(T - 1)$ rounds are non-binding and payoff-irrelevant. The moves chosen from M in the T-th round are binding commitments and determine the final outcome. The introduction of these rounds yields a *cheap play extension of* Γ, and is denoted Γ^T.

Let M^t denote the t-fold product of M for all $t \in \{1, \ldots, T\}$. For a given sequence of play, $m^T \in M^T$, $m(t)$ is used to denote the profile of moves chosen in the t-th round. An information set for player i at any $t > 1$ in the extensive-form game induced by Γ^T (referred to in the sequel, simply, as Γ^T) is characterized by a pair $(\theta_i, m^{t-1}) \in \Theta_i \times M^{t-1}$ and an information set for i at $t = 1$ is, obviously, $\theta_i \in \Theta_i$. Let $\mathcal{H}$ denote the

set of all information sets in Γ^T, with $\mathcal{H}_i$ identifying the subset of information sets that belong to player i. A strategy in Γ^T for player i is a function σ_i: $\mathcal{H}_i \to \Delta(M_i)$ and is given by a sequence $(s_i^t)_{t=1}^T$, where $s_i^1 \in S_i$, and for all $t > 1$, s_i^t: $\Theta_i \times M^{t-1} \to \Delta(M_i)$. Let Σ_i denote the strategy space for i in Γ^T. A *system of beliefs* is an n-tuple of functions β_i: $\Theta \times \bigcup_{t=1}^T M^t \to [0, 1]$ such that $\sum_{\theta_{-i} \in \Theta_{-i}} \beta_i(\theta, m^t) = 1$ for every $i \in N$, $\theta_i \in \Theta_i$ and $m^t \in \bigcup_{t'=1}^T M^{t'}$. β_i specifies the probability distribution that player i assigns to Θ_{-i}, given his/her information set. Let $\mathbb{B}$ denote the space of such systems of beliefs.

Given $\sigma = (s^t)_{t=1}^T$, let the function $g^*\sigma$: $\Theta \to \Delta(A)$ be defined by: for all $a \in A$ and $\theta \in \Theta$,

$$g^*\sigma(a \mid \theta) = \begin{cases} \sum_{m' \in \{m:\ g(m) = a\}} \tilde{\sigma}(m'(T) \mid \theta) & \text{if } \exists m \text{ such that } g(m) = a \\ 0 & \text{otherwise,} \end{cases}$$

where $\tilde{\sigma}(m'(T) \mid \theta)$ denotes the probability that m' will be played in the T-th round under the strategy profile σ in state θ.

Also, for any $\bar{m}_i \in M_i$, for any $t' \in \{1,\ldots, T\}$, if $m_i(t) = \bar{m}_i$ for all $t \in \{1,\ldots, t'\}$, then we write $m_i^{t'}$ as $[\bar{m}_i]$.

A *sequential equilibrium* of Γ^T is a pair $(\sigma, \beta) \in \Sigma \times \mathbb{B}$ that satisfies:

$$\forall i \in N,\ \forall h_i = (\theta_i, m^t) \in \mathcal{H}_i,\ \forall \sigma_i' \in \Sigma_i,\ \sigma' \equiv (\sigma_i', \sigma_{-i}),$$

$$\sum_{\theta_{-i} \in \Theta_{-i}} \sum_{a \in A} \beta_i(\theta, m^t) g^*\sigma(a \mid \theta) u_i(a, \theta) \geq \sum_{\theta_{-i} \in \Theta_{-i}} \sum_{a \in A} \beta_i(\theta, m^t) g^*\sigma'(a \mid \theta) u_i(a, \theta)$$

where β is compatible with the use of Bayes' Law whenever it is applicable.

$\mathcal{E}_{\Sigma\times\mathbb{B}}(\Gamma^T) \subseteq \Sigma \times \mathbb{B}$ denotes the set of sequential equilibria of Γ^T, with $\mathcal{E}(\Gamma^T) \equiv \{g^*\sigma: \exists\beta \in \mathbb{B}$ such that $(\sigma, \beta) \in \mathcal{E}_{\Sigma\times\mathbb{B}}(\Gamma^T)\}$ denoting the set of induced sequential equilibrium outcomes of Γ^T.

3. EQUILIBRIA WITH NO REGRET

In this paper, we will concentrate on the subset of sequential equilibria of Γ^T that are "regret-free". Such equilibria have the property that despite the changes in the beliefs of the players over the rounds of cheap play, the choices made in the early rounds remain optimal even when the time to make a binding commitment arrives.

Regret-free equilibria are natural focal points in games with cheap play. Cheap play is meant to capture the flavor of non-binding negotiation prior to the making of binding agreements. A player who makes proposals in the negotiation phase that would be regretted if they were payoff relevant would not be taken seriously by other players. "Serious" negotiation requires the use of strategies in the cheap play round that would also be optimal in the binding round. Thus, some form of stationarity in strategies over time seems to reflect this notion of no regret. There are many possible ways to formalize stationarity. Several are presented below.

A sequential equilibrium (σ, β) is *type I-stationary* if $\sigma = (s_i^t)_{t=1}^T$ satisfies: $\forall t \in \{2,\ldots, T\}$, $\forall i \in N$, $\forall\theta_i \in \Theta_i$, $\forall m^{t-1} \in M^{t-1}$ such that m^{t-1} arises with positive probability in the equilibrium, $s_i^t(\theta_i, m^{t-1}) = s_i^1(\theta_i)$.

A sequential equilibrium (σ, β) is *type II-stationary* if $\sigma = (s_i^t)_{t=1}^T$ satisfies: $\forall t \in \{2,\ldots, T\}$, $\forall i \in N$, $\forall\theta_i \in \Theta_i$, $\forall m^{t-1} \in M^{t-1}$, $s_i^t(\bar{m}_i \mid \theta_i,$

$m^{t-1}) = 1$, provided $m_i^{t-1} = [\bar{m}_i]$ for some $\bar{m}_i \in M_i$, and $s_i^1(\bar{m}_i \mid \theta_i) > 0$.

A sequential equilibrium (σ, β) is *type III-stationary* if $\sigma = (s_i^t)_{t=1}^T$ satisfies: $\forall t \in \{2,\ldots, T\}$, $\forall i \in N$, $\forall \theta_i \in \Theta_i$, $\forall m^{t-1} \in M^{t-1}$, $s_i^t(\theta_i, m^{t-1}) = s_i^1(\theta_i)$.

A sequential equilibrium (σ, β) is *type IV-stationary* if $\sigma = (s_i^t)_{t=1}^T$ satisfies: $\forall t \in \{2,\ldots, T\}$, $\forall i \in N$, $\forall \theta_i \in \Theta_i$, $\forall m^{t-1} \in M^{t-1}$, $s_i^t(m_i(t-1) \mid \theta_i, m^{t-1}) = 1$.

A sequential equilibrium (σ, β) is *type V-stationary* if $\sigma = (s_i^t)_{t=1}^T$ satisfies: $\forall t \in \{2,\ldots, T\}$, $\forall i \in N$, $\forall \theta_i \in \Theta_i$, $\forall m^{t-1} \in M^{t-1}$, $s_i^t(m_i(1) \mid \theta_i, m^{t-1}) = 1$.

A sequential equilibrium (σ, β) is *type VI-stationary* if $\sigma = (s_i^t)_{t=1}^T$ satisfies: $\forall t \in \{2,\ldots, T\}$, $\forall i \in N$, $\forall \theta_i \in \Theta_i$, $\forall m^{t-1} \in M^{t-1}$, $s_i^t(m_i(t') \mid \theta_i, m^{t-1}) = 1$, provided $t' < t$.

A sequential equilibrium (σ, β) is *type VII-stationary* if $\sigma = (s_i^t)_{t=1}^T$ satisfies: $\forall t \in \{2,\ldots, T\}$, $\forall i \in N$, $\forall \theta_i \in \Theta_i$, $\forall m^{t-1} \in M^{t-1}$, $s_i^t(\bar{m}_i \mid \theta_i, m^{t-1}) = 1$, provided $m_i^{t-1} = [\bar{m}_i]$ for some $\bar{m}_i \in M_i$.

There are fundamentally two categories of stationary strategies depending on the interpretation that one gives to mixed strategies. Type I stationarity characterizes the situation where the players choose a particular randomization over moves in the first period and choose the same randomization in every subsequent round, provided there are no deflections from the equilibrium path. This formulation implies that players choose a randomization that they will not regret regardless of what information is acquired during the remaining pre-play communication stages. This specification of no regret accepts the traditional viewpoint on mixed stategies which maintains that players actually randomize over alternative

actions.

The "modern" viewpoint, advocated by Harsanyi (1973), Aumann (1987), and others, is that mixed stategies are expressions of players' uncertainty about other players' actions. Thus, no player actually randomizes over actions, but for each player such a distribution represents other players' beliefs about his/her likely strategy choices. A no regret criterion that accepts this interpretation is captured by Type *II*-stationarity. It requires that in every round, the players choose the move they made in the first round, regardless of the information they receive in the interim, provided that (a) in the past they have chosen the first round move at every round and (b) the first round choice was on the equilibrium path. Note that mixtures (which are proxies for the players' beliefs) are considered only in the first round and subsequently only pure strategies appear since the uncertainty among the players about each other's moves is resolved.

Types *I* and *II* would seem to be the only natural notions of stationarity as long as our concern is only about no regret along the equilibrium path. However, a case can be made for insisting upon regret-free behavior off the equilibrium path as well. There are many reasonable alternatives that seem to capture such a notion. We describe five possibilities.

In Type *III*-stationarity, players are expected to choose a particular randomization over moves in the first period and choose the same randomization in every subsequent round, regardless of the move that is actually played in each round. Such a move may or may not occur with positve probability in equilibrium. In Type *IV*-stationarity, players choose to play the move that they had chosen in the previous round,

regardless of the information conveyed by earlier choices. In Type *V*-stationarity, in every round, the players choose the move they made in the first round, regardless of the information they receive in the interim. In Type *VI*-stationarity, at any round players choose one of the moves chosen in the past. Finally, in Type *VII*-stationarity, in every round of play, the players choose the move they made in the first round, regardless of the information they receive in the interim, provided in the past they have chosen the first round move at every round. Note that in all these cases, the move that is being repeated in subsequent rounds may or may not occur with positive proabability in equilibrium.

We could probably add many other specifications of behavior off the equilibrium path. We show that this is irrelevant in the next section. Other specifications would not differ from the ones presented here in any essential way.

4. RESULTS

Given the wealth of alternative no regret criteria, it is important to identify the one that is the most desirable from both practical and game-theoretic standpoints. In this section, we shall identify one such criterion and provide a complete characterization of the set of sequential equilibrium outcomes that arises under this restriction.

At the outset, we shall argue that Types *IV-VII*, which restrict behavior off the equilibrium path as well, are undesirable characterizations of regret-freeness. The restrictions that they impose on the set of sequential equilibria turn out to be too strong. This is the

message of the following proposition. Needless to say, Types *IV-VII* do not exhaust the set of possible variations on definitions of stationarity off the equilibrium path. However, the argument given below can be adapted to apply to any alternative definition which requires disequilibrium moves made in the cheap play rounds to be repeated in the binding round.

Theorem 1: *Generically, given* $T > 1$, $\mathcal{E}^{\nu}(\Gamma^{T}) = \varnothing$ *for* $\nu = IV, V, VI, VII$.

Proof: We shall give the argument for $T = 2$, in which case $\mathcal{E}^{\nu}(\Gamma^{T})$ is identical across $\nu = IV,\ldots, VII$. The arguments for arbitrary T are analogous.

Suppose θ is the state of the world and in round 1, m has been played. By Type ν-stationarity of the equilibrium strategy, ($\nu = IV,\ldots, VII$) m must be played again in round 2. By sequential rationality, m_i must be an optimal move against m_{-i} for all $i \in N$. Since this must hold for every $m \in M$, each $m_i \in M_i$ must be a best response to every $m_{-i} \in M_{-i}$ for all $i \in N$, i.e. each $m \in M$ must yield the same utility to any player i. Hence, the game in any given state of the world must be trivial and a slight perturbation of payoffs would destroy the equilibrium. ∎

Given that the remaining choices are Types *I-III*, the main objective of this paper is to provide a strong case for selecting Type *II* stationarity as the appropriate formulation of no regret. This is based on both game-theoretic and practical arguments.

From a game-theoretic standpoint, we argue against Types *I* and *III* as appropriate formulations of regret-freeness. The question reduces to the preferred interpretation of mixed strategies. The traditional view of randomization that these types embody is that players would base important

decisions on the flip of a coin or some other randomization device. This, as Aumann (1989) puts it, is difficult to swallow. Following Savage (1954), we take the view (espoused by Harsanyi (1973), Aumann (1987) and others) that randomness is not a physical act; a mixed strategy for player i is simply an expression of the ignorance that other players have about i's decisions.

Of course, a case can be made for the traditional view. For instance, poker players or auditors and quality inspectors are known to choose randomizations to avoid becoming predictable. In general, such justifications would not apply and, moreover, run into conceptual difficulties. Thus, the modern view of mixed strategies supports the choice of Type *II* stationarity.

The second argument in favor of Type *II* stationarity is that the resulting set of equilibrium outcomes has a simple characterization. The practical implications of the characterization is that we need not be concerned with the messy extended cheap play game. Instead it is sufficient to focus on a subset of the Bayesian equilibria of the underlying static game.

This characterization is in terms of a simple condition on the set of Bayesian equilibria of the underlying static game Γ and is called *posterior implementability*. It was introduced by Green and Laffont (1987). This notion is defined below. For all $i \in N$ and $s \in S$, define the function π_i^s: $\Theta \times M \to [0, 1]$ by

$$\pi_i^s(\theta, m) = \frac{s_{-i}(m_{-i} \mid \theta_{-i})\pi_i(\theta)}{\sum_{\theta'_{-i} \in \Theta_{-i}} s_{-i}(m_{-i} \mid \theta'_{-i})\pi_i(\theta'_{-i}, \theta_i)}$$

A Bayesian equilibrium s is *posterior implementable* if

$$\forall \theta \in \Theta,\ \forall m \in M \text{ such that } s(m \mid \theta) > 0,\ \forall i \in N,\ \forall m_i' \in M_i,$$

$$\sum_{\theta_{-i} \in \Theta_{-i}} \pi_i^s(\theta, m) u_i(g(m), \theta) \geq \sum_{\theta_{-i} \in \Theta_{-i}} \pi_i^s(\theta, m) u_i(g(m_i', m_{-i}), \theta)$$

$\mathcal{E}_S^{PI}(\Gamma) \subseteq S$ denotes the set of posterior implementable Bayesian equilibria of Γ, with $\mathcal{E}^{PI}(\Gamma) \equiv \{g^*s\colon s \in \mathcal{E}_S^{PI}(\Gamma)\}$ denoting the set of Bayesian equilibrium outcomes of Γ satisfying posterior implementability.

Two theorems follow. The first shows that Type *II* stationarity is characterized by the posterior implementability criterion. The second shows that the Types *I* and *III* have no logical relationship with posterior implementability.

Theorem 2: $\mathcal{E}^{II}(\Gamma^T) = \mathcal{E}^{PI}(\Gamma)$ *for all* $T > 1$.

Proof: The theorem is a consequence of the following lemmata.

In the proofs below, if $m^{t-1} \in M^{t-1}$ and $s \in S$ are such that $m^{t-1} = [m]$ for some $m \in M$ with $s(m \mid \theta) > 0$, then we shall say that (m^{t-1}, m, s) satisfy Property A.

Lemma 1: $\mathcal{E}^{PI}(\Gamma) \subseteq \mathcal{E}^{II}(\Gamma^T)$.

Proof of Lemma 1: Choose $s \in \mathcal{E}^{PI}(\Gamma)$. Also, choose $\sigma \in \Sigma$ such that $\sigma = (s^t)_{t=1}^T$ with

(a) $s^1 = s$,

(b) for all $t \in \{2, \ldots, T\}$, for all $\theta \in \Theta$, for all $m^{t-1} \in M^{t-1}$, $s^t(m \mid \theta, m^{t-1}) = 1$ if (m^{t-1}, m, s) satisfy Property A, and

(c) for all $\theta \in \Theta$, and all $m^{T-1} \in M^{T-1}$, such that (m^{T-1}, m, s) does

not satisfy Property A for any m, $s^T(\cdot \mid \theta, m^{T-1}) = s^1(\cdot \mid \theta)$.

Define $\beta \in \mathbb{B}$ such that it satisfies:

for all $i \in N$, for all $\theta \in \Theta$, for all $m^{t-1} \in M^{t-1}$,

if there exists $m \in M$ such that (m^{t-1}, m, s) satisfy Property A,

$$\beta_i(\theta, m^{t-1}) = \frac{s^1_{-i}(m_{-i} \mid \theta_{-i})\pi_i(\theta)}{\sum_{\theta'_{-i} \in \Theta_{-i}} s^1_{-i}(m_{-i} \mid \theta'_{-i})\pi_i(\theta'_{-i}, \theta_i)}$$

otherwise,

$$\beta_i(\theta, m^{t-1}) = \pi_i(\theta).$$

In words, along the equilibrium path, the players use strategy profile s in period 1, and subsequently repeat the move chosen in the first period with probability one. Correspondingly, beliefs are updated by an application of Bayes' Law. Out of equilibrum, players' beliefs are identical to the original posteriors and players use the strategy profile s.

By definition of $\mathcal{E}^{PI}(\Gamma)$, for all $m^{t-1} \in M^{t-1}$ and $m \in M$ such that (m^{t-1}, m, s) satisfy Property A, for all $i \in N$, for all $\theta \in \Theta$, for all $\tilde{m}_i \in M_i$,

$$\sum_{\theta_{-i} \in \Theta_{-i}} \beta_i(\theta, [m])u_i(g(m), \theta) \geq \sum_{\theta_{-i} \in \Theta_{-i}} \beta_i(\theta, [m])u_i(g(m_{-i}, \tilde{m}_i), \theta).$$

Moreover, since $s \in \mathcal{E}_S(\Gamma)$, for all $m^{t-1} \in M^{t-1}$ such that $m^{t-1} \neq [m]$ for any $m \in M$ with $s(m \mid \theta) > 0$, we have for all $i \in N$, for all $\theta \in \Theta$, for all $\tilde{m}_i \in M_i$,

$$\sum_{m_{-1}\in M_{-1}} \sum_{\theta_{-1}\in\Theta_{-1}} s^1_{-1}(m_{-1} \mid \theta_{-1})\beta_1(\theta, m^{t-1})u_1(g(m), \theta) \geq$$

$$\sum_{m_{-i}\in M_{-1}} \sum_{\theta_{-1}\in\Theta_{-i}} s^1_{-1}(m_{-1} \mid \theta_{-1})\beta_1(\theta, m^{t-1})u_1(g(m_{-1}, \tilde{m}_i), \theta).$$

By definition, $(\sigma, \beta) \in \mathcal{E}_{\Sigma\times\mathbb{B}}(\Gamma^T)$ and satisfies Type *II*-stationarity. By construction, $g^*s = g^*\sigma$. Thus, $\mathcal{E}^{PI}(\Gamma) \subseteq \mathcal{E}^{II}(\Gamma^T)$. ∎

Lemma 2: $\mathcal{E}^{II}(\Gamma^T) \subseteq \mathcal{E}^{PI}(\Gamma)$.

Proof of Lemma 2: Choose $(\sigma, \beta) \in \mathcal{E}^{II}_{\Sigma\times\mathbb{B}}(\Gamma^T)$, with $\sigma = (s^t)^T_{t=1}$. By sequential rationality, and Type *II*-stationarity, for all $i \in N$, for all $\theta_i \in \Theta_i$, for all $m \in M$, such that $s^1(m \mid \theta) > 0$, for all $\tilde{m}_i \in M_i$, [1] must hold:

$$\sum_{\theta_{-1}\in\Theta_{-i}} \beta_1(\theta, [m])u_1(g(m), \theta) \geq \sum_{\theta_{-1}\in\Theta_{-i}} \beta_1(\theta, [m])u_1(g(m_{-1}, \tilde{m}_1), \theta). \qquad [1]$$

But by Type *II*-stationarity,

$$\beta_1(\theta, [m]) = \frac{s^1_{-i}(m_{-i} \mid \theta_{-i})\pi_i(\theta)}{\sum_{\theta'_{-1}\in\Theta_{-1}} s^1_{-i}(m_{-i} \mid \theta'_{-i})\pi_i(\theta'_{-i}, \theta_i)}$$

$$= [s^1_{-i}(m_{-i} \mid \theta_{-1})\pi_i(\theta)]K(m_{-1}),$$

where $K(m_{-i})$ is a constant term.

Hence, by [1] the posterior implementability criterion is met. Next, we need to check that $\mathcal{E}^{II}(\Gamma^T) \subseteq \mathcal{E}(\Gamma)$.

By substituting in [1] and multiplying through by $[K(m_{-1})]^{-1}$, for all $i \in N$, for all $\theta_i \in \Theta_i$, for all $m \in M$ such that $s^1(m \mid \theta) > 0$, for all $\tilde{m}_i \in M_i$, [2] must hold:

$$\sum_{\theta_{-i} \in \Theta_{-i}} s^1_{-i}(m_{-i} \mid \theta_{-i}) \pi_i(\theta) u_i(g(m), \theta) \geq$$

$$\sum_{\theta_{-i} \in \Theta_{-i}} s^1_{-i}(m_{-i} \mid \theta_{-i}) \pi_i(\theta) u_i(g(m_{-i}, \tilde{m}_i), \theta). \qquad [2]$$

Since [2] is true for each $m_{-i} \in M_{-i}$, such that $s^1_{-i}(m_{-i} \mid \theta_{-i}) > 0$, for all $i \in N$, for all $\theta_i \in \Theta_i$, for all $m_i \in M_i$ such that $s^1_i(m_i \mid \theta_i) > 0$, and for all $\tilde{m}_i \in M_i$, [3] must hold:

$$\sum_{m_{-i} \in M_{-i}} \sum_{\theta_{-i} \in \Theta_{-i}} s^1_{-i}(m_{-i} \mid \theta_{-i}) \pi_i(\theta) u_i(g(m), \theta) \geq$$

$$\sum_{m_{-i} \in M_{-i}} \sum_{\theta_{-i} \in \Theta_{-i}} s^1_{-i}(m_{-i} \mid \theta_{-i}) \pi_i(\theta) u_i(g(m_{-i}, \tilde{m}_i), \theta). \qquad [3]$$

Thus, $s^1 \in \mathcal{E}_S(\Gamma)$. By Type *II*-stationarity, $g^*\sigma = g^*s^1$. Hence $\mathcal{E}^{II}(\Gamma^T) \subseteq \mathcal{E}^{PI}(\Gamma)$. ■

Theorem 3: *There is no logical relationship either between* $\mathcal{E}^I(\Gamma^T)$ *and* $\mathcal{E}^{PI}(\Gamma)$ *or between* $\mathcal{E}^{III}(\Gamma^T)$ *and* $\mathcal{E}^{PI}(\Gamma)$ *for any* $T > 1$.

Proof: The following examples are proof of the proposition above. The first example shows that $\mathcal{E}^I(\Gamma^T)$ does not contain $\mathcal{E}^{PI}(\Gamma)$, and the second example shows that the reverse containment may not be expected either. By construction, in the examples below, every move is assigned positive probability in equilibrium. Hence, the argument would hold for the case of $\mathcal{E}^{III}(\Gamma^T)$ as well. Any non-genericity in the payoffs is purely for the purposes of keeping the examples as simple as possible and is not critical for the arguments.

Example 1:

Consider the following game, Γ.

[INSERT FIGURE 1 HERE]

The players are labelled 1 and 2 and they choose from the sets $\{U, D\}$ and $\{L, R\}$ respectively. Player 2 has private information; her set of possible private observations is $\{\theta_2, \theta_2'\}$. Player 1 is uninformed and has a single possible observation. Player 1's posterior distribution on Θ_2 is $\pi_1(\theta_2) = \pi_1(\theta_2') = \frac{1}{2}$.

Trivially, player 2 is indifferent among her strategies. Consider the following Bayesian equilibrium s:

$$s_1(U \mid \theta_1) = \frac{1}{2}.$$
$$s_2(L \mid \theta_2) = \frac{2}{3}; \quad s_2(L \mid \theta_2') = \frac{1}{3}.$$

Both L and R are assigned positive probability by the equilibrium. Conditional upon the observation of L, the probability that player 1 assigns to θ_2 is $\frac{2}{3}$ and in the event that R is observed, the probability assigned to θ_2 is $\frac{1}{3}$. Upon observing either one of player 2's moves, the payoff to player 1 from choosing either U or D is $\frac{10}{3}$. Hence, the Bayesian equilibrium given above is also posterior implementable.

The equilibrium s induces the following distribution on the outcome space in each state.

[INSERT FIGURE 2]

Suppose that there exists a Type I-stationary strategy in the game Γ^T, denoted $\sigma = (\hat{s}^t)_{t=1}^T$, that also yields the same distribution on outcomes. By Type I-stationarity, $\hat{s}^1$ must satisfy:

$$\begin{bmatrix} \hat{s}_1^1(U|\ \theta_1) \\ \hat{s}_1^1(D|\ \theta_1) \end{bmatrix} \begin{bmatrix} \hat{s}_2^1(L|\ \theta_2) & \hat{s}_2^1(R|\ \theta_2) \end{bmatrix} = \begin{bmatrix} \frac{1}{3} & \frac{1}{6} \\ \frac{1}{3} & \frac{1}{6} \end{bmatrix}$$

$$\begin{bmatrix} \hat{s}_1^1(U|\ \theta_1) \\ \hat{s}_1^1(D|\ \theta_1) \end{bmatrix} \begin{bmatrix} \hat{s}_2^1(L|\ \theta_2') & \hat{s}_2^1(R|\ \theta_2') \end{bmatrix} = \begin{bmatrix} \frac{1}{6} & \frac{1}{3} \\ \frac{1}{6} & \frac{1}{3} \end{bmatrix}$$

The unique solution to the equations above is $\hat{s}^1 = s$. However, σ is not an equilibrium strategy. Suppose that player 2 chooses to play L in the first round. Player 1's payoff from playing U in the second round is:

$$\tfrac{2}{3}(\tfrac{2}{3}(5) + \tfrac{1}{3}(2)) + \tfrac{1}{3}(\tfrac{1}{3}(0) + \tfrac{2}{3}(4)) = \tfrac{32}{9},$$

which is strictly greater than the payoff from playing D in the second round, given by:

$$\tfrac{2}{3}(\tfrac{2}{3}(4) + \tfrac{1}{3}(0)) + \tfrac{1}{3}(\tfrac{1}{3}(2) + \tfrac{2}{3}(5)) = \tfrac{28}{9}.$$

The strategy σ_1 is not a best response to σ_2 for the case $T = 2$. It can be checked that the same conclusion would be obtained for arbitrary T (> 1). Thus, $\mathcal{E}^{PI}(\Gamma)$ is not a subset of $\mathcal{E}^I(\Gamma^T)$.

Example 2:

Consider the following game, Γ.

[INSERT FIGURE 3]

The players are labelled 1 and 2 and they choose from the sets $\{U, D\}$ and $\{L, R\}$ respectively. Player 1 has private information; his set of possible private observations is $\{\theta_1, \theta_1'\}$. Player 2 is uninformed. Player 2's posterior distribution on Θ_1 is $\pi_2(\theta_1) = \pi_2(\theta_1') = \frac{1}{2}$.

Trivially, player 2 is indifferent among her strategies. Consider the following Bayesian equilibrium s:

$$s_1(U \mid \theta_1) = \frac{2}{3}; \quad s_1(U \mid \theta_1') = \frac{1}{3}.$$

$$s_2(L \mid \theta_2) = \frac{1}{2}.$$

It is easily checked that $\sigma = (s^t = s)_{t=1}^{T}$ is a sequential equilibrium strategy profile (which satisfies Type I-stationarity, by definition) as well. For any $T > 1$, $T - 1$ rounds of cheap play do not invalidate the best-response property of either player's strategy.

The equilibrium σ induces the following distribution on the outcome space in each state.

[INSERT FIGURE 4]

Let $\hat{s} \in S$ be a strategy profile in the game Γ. If $\hat{s}$ must yield the distribution over outcomes given above, it must satisfy:

$$\begin{bmatrix} \hat{s}_1(U \mid \theta_1) \\ \hat{s}_1(D \mid \theta_1) \end{bmatrix} \begin{bmatrix} \hat{s}_2(L \mid \theta_2) & \hat{s}_2(R \mid \theta_2) \end{bmatrix} = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} \\ \frac{1}{6} & \frac{1}{6} \end{bmatrix}$$

$$\begin{bmatrix} \hat{s}_1(U \mid \theta_1') \\ \hat{s}_1(D \mid \theta_1') \end{bmatrix} \begin{bmatrix} \hat{s}_2(L \mid \theta_2) & \hat{s}_2(R \mid \theta_2) \end{bmatrix} = \begin{bmatrix} \frac{1}{6} & \frac{1}{6} \\ \frac{1}{3} & \frac{1}{3} \end{bmatrix}$$

The unique solution to the equations above is $\hat{s} = s$. However, s is not posterior implementable. Suppose that player 1's observation is θ_1 and he observes player 2's play of R. Player 1's payoff from playing U clearly dominates the payoff from playing D, given this observation.

Thus, $\mathcal{E}^I(\Gamma^T)$ is not a subset of $\mathcal{E}^{PI}(\Gamma)$. ■

5. CONCLUDING REMARKS

In this paper, we study pre-play communication via the announcements of moves that are non-binding and payoff irrelevant. We argue that sequential equilibria of the induced extended game that satisfy a criterion of no regret are focal. We find that there are several ways to formalize the notion of no regret and present an argument for choosing one of them. Our main result gives a characterization of the set of regret-free sequential equilibrium outcomes in terms of a posterior implementability restriction on the equilibria of the underlying static game. This notion of regret-freeness is consistent with the modern interpretation of strategies proposed by Harsanyi (1973) and Aumann (1987). Finally, we show that notions of regret-freeness consistent with the traditional view of mixed strategies have no logical connection with posterior implementability.

In Chakravorti (1990), a special case of the analysis given earlier is applied. In that paper, the equivalence between the stationary pure

strategy Bayesian equilibria of the extended game and pure strategy posterior implementable equilibria of the static game is considered. Note that all of the definitions of stationarity given in this paper are equivalent when attention is restricted to pure strategies and to regret-freeness only in parts of the game that are on the equilibrium path.

We close with two observations.

First, a common theme underlying the various stationarity notions considered in this paper is that there is no regret at every round on the equilibrium path -- "uniform" no regret. One could, of course, conceive of other notions that convey regret-freeness in "cycles". For example, for the first three rounds, players experiment and subsequently repeat the moves made three periods earlier; or the players oscillate between two equilibria in alternate rounds, etc. We leave such issues open for future research. Our intention in this paper is to focus entirely on "uniform" no regret without exploring the Pandora's box of variations on cyclical definitions.

Second, we note that the equivalence result given in the previous section may continue to hold even in the case of games without completely cheap play. In Chakravorti and Conley (1991), a theory of play with random deadlines for submitting a binding action is introduced. In the current context, this would mean that there is some positive probability (less than one) with which the play in the preliminary rounds is not cheap, i.e. the game could be terminated and any moves made would be used to compute payoffs. This rules out "babbling" sequential equilibria of the extended game and gives players an incentive to "negotiate in good faith." We speculate that the result of this paper would be unaffected by the presence of such a possibility.

References

AUMANN, R. (1976): "Agreeing to Disagree," *Annals of Stat.* **4**:1236-1239.

AUMANN, R. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica* **55**:1-18.

AUMANN, R. (1989): "Game Theory" in *The New Palgrave,* J. Eatwell, M. Milgate and P. Newman (eds.), New York, Norton.

CHAKRAVORTI, B. (1990): "Efficiency and Mechanisms with No Regret," forthcoming in *Int. Econ. Rev.* (1992).

CHAKRAVORTI, B. AND J. CONLEY (1991): "On Games with Random Deadlines," forthcoming.

FARRELL, J. (1982): "Communication in Games," mimeo, MIT.

GREEN, J. AND J.-J. LAFFONT (1987): "Posterior Implementability in a Two-Person Decision Problem," *Econometrica*, **55**:69-94.

HARSANYI, J. (1973): "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points," *Int. J. Game Theory* 2:1-23.

SAVAGE, L. (1954): *The Foundations of Statistics*, New York. John Wiley.

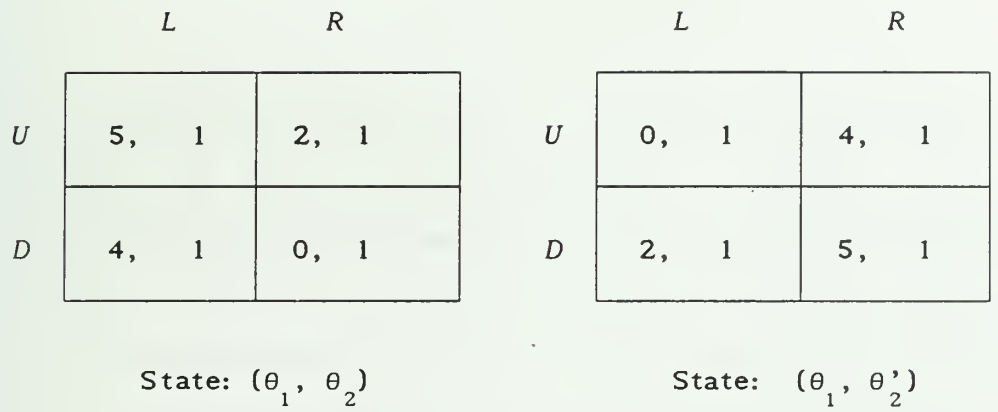

FIGURE 1.

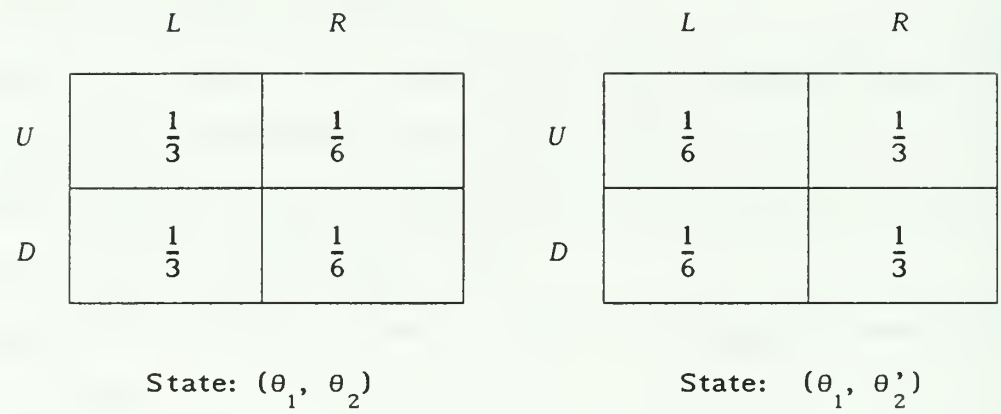

FIGURE 2.

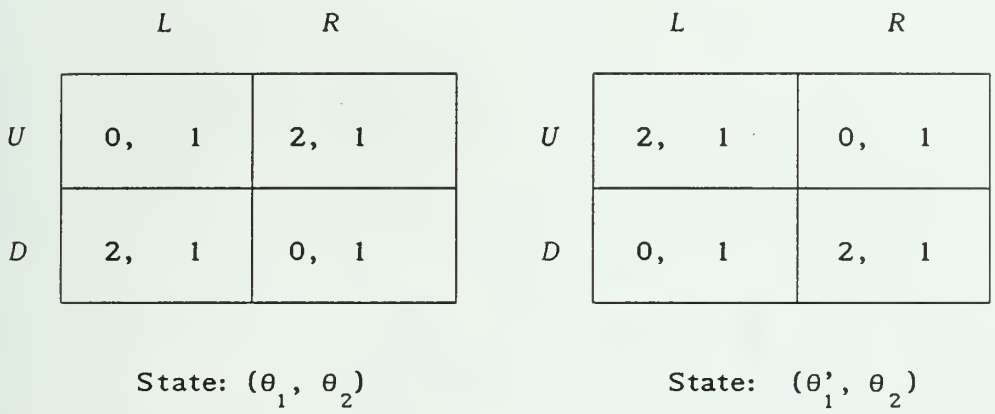

FIGURE 3.

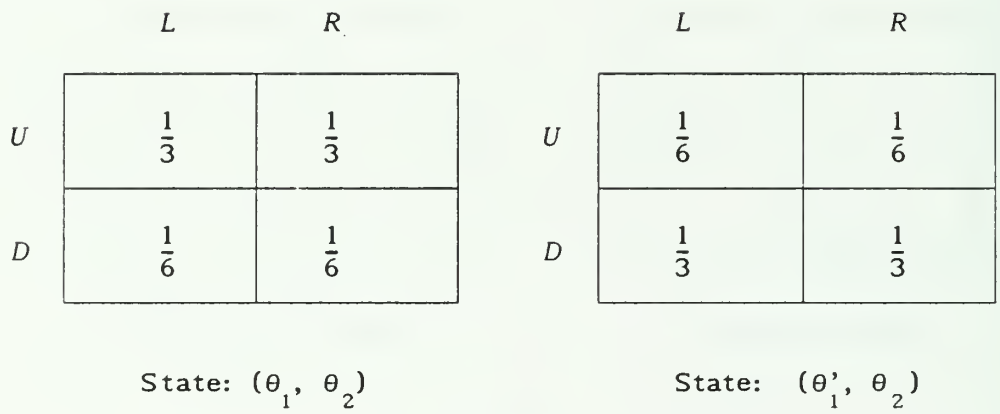

FIGURE 4.

HECKMAN
BINDERY INC.
JUN 95
Bound-To-Please® N. MANCHESTER,
INDIANA 46962

UNIVERSITY OF ILLINOIS-URBANA
3 0112 060295901